

Faraday synthesis

The synergy of aperture and rotation measure synthesis

M.R. Bell¹ and T.A. Enßlin¹

Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany e-mail: mrbell@mpa-garching.mpg.de

Received ??? / Accepted ???

Abstract

We introduce a new technique for imaging the polarized radio sky using interferometric data. The new approach, which we call Faraday synthesis, combines aperture and rotation measure synthesis imaging and deconvolution into a single algorithm. This has several inherent advantages over the traditional two-step technique, including improved sky plane resolution, fidelity, and dynamic range. In addition, the direct visibility- to Faraday-space imaging approach is a more sound foundation on which to build more sophisticated deconvolution or inference algorithms. For testing purposes, we have implemented a basic Faraday synthesis imaging software package including a three-dimensional CLEAN deconvolution algorithm. We compare the results of this new technique to those of the traditional approach using mock data. We find many artifacts in the images made using the traditional approach that are not present in the Faraday synthesis results. In all, we achieve a higher spatial resolution, an improvement in dynamic range of about 20%, and a more accurate reconstruction of low signal to noise source fluxes when using the Faraday synthesis technique.

Key words. Polarization - Techniques: interferometric - Techniques: polarimetric - Methods: data analysis - Magnetic Fields

1. Introduction

Rotation measure (RM) synthesis, introduced by Brentjens & de Bruyn (2005), is a technique that makes use of the Faraday effect to improve the sensitivity of polarimetric observations by combining data over wide ranges in frequency. RM synthesis allows for the separation of polarized sources along the line of sight (LOS) by decomposing the observed polarized emission into parts originating from different Faraday depths (in the simplest case, Faraday depth is equivalent to RM), allowing one to generate a 3D representation of the polarized sky. While the Faraday depth axis cannot be mapped to a physical depth, the Faraday depth information can be of significant scientific value since the Faraday depth traces the projected strength and orientation of magnetic fields along the LOS. A more detailed introduction to RM synthesis is provided later.

The RM synthesis technique has only recently become viable, owing to the availability of broadband receivers in the next generation of radio telescopes such as the Expanded Very Large Array (EVLA), the upgraded Westerbork Synthesis Radio Telescope (WSRT), and the pathfinder projects leading to the Square Kilometer Array (SKA) such as the Low-Frequency Array (LOFAR). Recently, there have been many successful applications of RM synthesis, and interest is rapidly increasing as new radio telescopes are being commissioned. Applications have included studies of the diffuse polarized emission in the Perseus field (de Bruyn & Brentjens 2005; Brentjens 2011) and the Abell 2255 field (Pizzo et al. 2010), analysis of the polarized emission in nearby galaxies in the WSRT-SINGS survey (Heald et al. 2009), and the detection of a shell of compressed magnetic fields surrounding a local HI bubble (Wolleben et al. 2010). The RM synthesis technique will play a critical role in several upcoming polarization surveys,

e.g. POSSUM (Gaensler et al. 2010), GMIMS (Wolleben et al. 2009), and future surveys with LOFAR. RM synthesis is very useful for studying magnetism. For instance, Bell et al. (2011) have shown that prominent asymmetric features in RM synthesis images known as Faraday caustics, which are related to LOS magnetic field reversals, can be used to study the structural and statistical properties of magnetic fields.

In addition to the considerable interest in applications, there has been recent interest in improving RM synthesis imaging techniques. The RMCLEAN deconvolution algorithm was introduced by Heald et al. (2009) and was quickly adopted, owing to its simplicity and similarity to techniques used in aperture synthesis imaging. Frick et al. (2010) have proposed a wavelet-based RM synthesis technique. With this approach, one obtains a decomposition of the size scale of structures in addition to their Faraday depth location. There has also been growing interest in applying compressed sensing to RM synthesis (Li et al. 2011; Andreut et al. 2011). Compressed sensing is a method of reconstructing signals that are sparse in some set of basis functions. If the signal is sufficiently sparse, it can be reconstructed using fewer measurements than indicated by the Nyquist-Shannon sampling theorem. Since compressed sensing techniques make use of wavelet bases, but are implemented in a noise-aware fashion, they can be expected to be superior to pure wavelet based methods.

These new imaging techniques have thus far focused on the problem of one-dimensional (1D) reconstruction. However, the product of RM synthesis is a function not only of Faraday depth, but also of position on the sky, making its reconstruction an inherently three-dimensional (3D) problem. In many cases, RM synthesis is performed on sky brightness images that have been produced from radio interferometric data. Observations performed with a radio interferometer sample the aperture plane

rather than the image plane, and this is done at many different frequencies. This data space is the 3D Fourier space representation of the polarized sky brightness as a function of Faraday depth. Imaging algorithms should ideally make use of the entire data space to inform the reconstruction at each pixel, since the information about the sky brightness at each pixel is spread throughout Fourier space. However, the current approach is to perform the imaging in a piecewise fashion, first reconstructing the 2D sky plane images at each frequency before doing 1D RM synthesis imaging along each LOS. Therefore, each step of the traditional imaging approach is done with a limited subset of the data, which will reduce the overall sensitivity and degrade fidelity in the final image.

In this paper, we introduce a new technique for imaging the Faraday spectrum directly from radio interferometric data. We call this technique *Faraday synthesis*. Using this approach, one images the polarized emission as a function of sky position and Faraday depth from the visibility data itself, rather than using the traditional piecewise prescription. Faraday synthesis is a natural extension of the aperture synthesis plus RM synthesis techniques that provides improvements in image fidelity and sensitivity.

We note that a similar technique was briefly discussed in Pen et al. (2009), although it was not considered in any detail, nor was it compared to the traditional approach to RM synthesis imaging. Furthermore, deconvolution was not considered.

With the advent of RM synthesis the concept of rotation measure, defined to be the amount that the observed polarization angle changes as a function of frequency, has become somewhat outdated. With RM synthesis, one does not measure RMs, but instead reconstructs the polarized intensity as a function of Faraday depth. In the simplest case, where a single, discrete source of polarized emission is positioned behind a Faraday rotating medium, the RM is equal to the Faraday depth. In all other cases this is not true. In general, RM cannot be used as a proxy for Faraday depth, and the full distribution of polarized brightness as a function of Faraday depth is the most appropriate quantity to study. Therefore, we avoid use of the term RM to describe this new method, and instead call it Faraday synthesis. Throughout the remainder of this paper, RM synthesis will refer to the LOS imaging method developed by Brentjens & de Bruyn (2005). The traditional practical approach of first imaging individual frequencies using 2D aperture synthesis techniques and then reconstructing the LOS brightness distribution on a pixel-by-pixel basis will be referred to as 2+1D imaging, in contrast to Faraday synthesis, which we will often refer to as 3D imaging.

In Sec. 2 we briefly review the theories of aperture and RM synthesis imaging. In Sec. 3 we introduce the Faraday synthesis imaging technique. In Sec. 4 we describe the proof of concept software that we have implemented, and in Sec. 5 we compare test results obtained by imaging mock data using both the 3D and 2+1D techniques. We conclude in Sec. 6 with a summary and discussion of our results.

2. Synthesis imaging

In this section we review the fundamentals of aperture and RM synthesis imaging before showing how they can be performed simultaneously in Faraday synthesis imaging. Reviews are included here for completeness and to highlight the assumptions that are typically made and the limitations that result.

2.1. Aperture synthesis

We can not possibly provide a complete review of the theory of aperture synthesis. We only wish to review those aspects that are most relevant to the current work. For a comprehensive treatment, the reader is referred to Thompson et al. (2001).

With an interferometer, one measures not the sky brightness directly, but rather a collection of discrete samplings of the aperture plane. These samples are complex quantities, typically referred to as visibilities, denoted as V . The visibilities are the correlated voltage output of pairs of antennas. For a narrow-band observation, they are related to the sky brightness distribution, I , by

$$V(u, v, w, \lambda) = \iint_{-\infty}^{\infty} \frac{dl dm}{\sqrt{1-l^2-m^2}} I(l, m, \lambda) e^{-2\pi i [ul+vm+w(\sqrt{1-l^2-m^2}-1)]}. \quad (1)$$

The coordinates (u, v, w) are spatial frequency coordinates, or the distance between pairs of antennas, measured in numbers of wavelengths. The coordinate u measures the distance in the cardinal North-South direction, while v measures the distance in the East-West direction. The coordinate w points in the direction of the phase reference position on the sky. The (l, m) coordinates are direction cosines relative to the (u, v) coordinates. The wavelength at which the visibilities are measured is given by λ .

This relationship can be simplified to a two-dimensional (2D) Fourier transformation in two circumstances. The first is in the case of an East-West oriented array such as the WSRT. In this case, the telescopes move through a plane such that $w = 0$ as the Earth rotates. The second case is when only a small patch of the sky is being imaged, such that $w(\sqrt{1-l^2-m^2}-1) \approx 0$. For the time being, we assume that we are looking at a small patch of the sky and henceforth neglect this w -term.

A radio telescope is not equally sensitive to the entire sky. The sky brightness distribution is attenuated by the antenna power pattern, A , which is commonly referred to as the primary beam. Including this effect, the visibilities are related to the sky brightness distribution via the relationship

$$V'(u, v, \lambda) = \iint_{-\infty}^{\infty} dl dm A(l, m, \lambda) I(l, m, \lambda) e^{-2\pi i (ul+vm)}. \quad (2)$$

In reality, as mentioned above, only discrete locations in the aperture plane are sampled. The measured visibilities, \widehat{V} , are related to the true visibilities by

$$\widehat{V}(u, v, \lambda) = S(u, v, \lambda) V'(u, v, \lambda). \quad (3)$$

The sampling function S can be represented as

$$S(u, v, \lambda) = W(u, v, \lambda) \sum_i \delta\left(u - \frac{\mathbf{b}_i \cdot \hat{\mathbf{x}}}{\lambda}\right) \delta\left(v - \frac{\mathbf{b}_i \cdot \hat{\mathbf{y}}}{\lambda}\right) \delta(\lambda - \lambda_i) \quad (4)$$

where $\mathbf{b} = b_x \hat{\mathbf{x}} + b_y \hat{\mathbf{y}}$ is the distance between two antennas, known as the baseline length, and the unit vectors $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ point toward the North and East, respectively. The function W allows for the inclusion of weighting factors, e.g. by the inverse of the noise.

The i subscript is an index over the list of discrete values of \mathbf{b} and λ for which measurements have been made.

To recover the sky brightness distribution from visibility data, one must invert Eqs. 2 and 3. Due to the sampling function and the presence of noise, it is not possible to solve for I uniquely. The inverse Fourier transform of \hat{V} does not give the sky brightness distribution alone, but rather

$$\begin{aligned} I_D(l, m, \lambda) &= \iint_{-\infty}^{\infty} du dv \hat{V}(u, v, \lambda) e^{2\pi i(ul+vm)} \\ &= \iint_{-\infty}^{\infty} du dv S(u, v, \lambda) V'(u, v, \lambda) e^{2\pi i(ul+vm)} \\ &= B_{\text{sky}}(l, m, \lambda) *_{lm} [A(l, m, \lambda) I(l, m, \lambda)], \end{aligned} \quad (5)$$

where $*_{lm}$ denotes convolution in the l and m plane. The image I_D is commonly referred to as the *dirty image*. The *dirty beam*, B_{sky} , is

$$B_{\text{sky}} = \iint_{-\infty}^{\infty} du dv S(u, v, \lambda) e^{2\pi i(ul+vm)}. \quad (6)$$

Due to the extended structure of the dirty beam, the sky image obtained from Eq. 5 has a limited dynamic range and an unphysical appearance. Deconvolution is required to recover a reasonable approximation of I and to detect weaker features that are obscured by artifacts associated with the bright sources. By far, the most commonly used deconvolution algorithm is the CLEAN algorithm, first introduced by Högbom (1974) and later improved by Clark (1980) as well as Schwab (1984). We describe the CLEAN algorithm later when discussing the 3D implementation that has been included in our proof-of-concept software.

In the simplest case, when A is independent of baseline and time, the effect of the primary beam can be removed by simply dividing by a known beam pattern. In doing so, the flux scale will be normalized across the sky and the noise level will increase as a function of distance from the pointing center. In general, however, A (and other so-called direction dependent effects) can depend on both baseline and time, and in this case one must use something like the A-projection algorithm described by Bhatnagar et al. (2008).

2.2. RM synthesis

Faraday rotation is a birefringence effect where the plane of polarization of a plane-polarized wave is rotated as it passes through a magneto-ionic medium. The right- and left-circularly polarized components of the plane wave propagate at different speeds through the medium causing a relative phase shift, and therefore a rotation of the polarization plane. The amount of rotation incurred is given by

$$\chi(\lambda^2) = \chi_0 + \phi \lambda^2, \quad (7)$$

where χ is the observed position angle at wavelength λ , and χ_0 is the position angle at the source of emission. The quantity ϕ , known as the *Faraday depth*, is given by

$$\phi(z) = (0.81 \text{ rad/m}^2) \int_0^z \left(\frac{dz'}{\text{pc}} \right) \left[\frac{n_e(z')}{\text{cm}^{-3}} \right] \left[\frac{B_z(z')}{\mu\text{G}} \right], \quad (8)$$

where B_z is the LOS component of the magnetic field, and z is the distance along the LOS. The number density n_e includes both thermal electrons and positrons, which are given as negative and positive values, respectively. Faraday depth is distinct from RM, which we define following Burn (1966) to be

$$RM(\lambda^2) \equiv \frac{\partial \chi(\lambda^2)}{\partial \lambda^2}. \quad (9)$$

In the simplest case, when a single point source of polarized emission sits behind a Faraday rotating medium, the RM measured for the source is equal to ϕ . In all other cases, these quantities differ. In the general case of mixed Faraday rotating and synchrotron emitting media, the observed polarized intensity originates from a range of Faraday depths, and the RM varies as a function of λ^2 . The full polarized intensity as a function of Faraday depth, as obtained using RM synthesis, is required for an opportunity to study the intrinsic properties of the various sources along the LOS.

The polarized intensity as a function of sky position and wavelength is a complex quantity that is usually defined in terms of the Stokes parameters to be

$$P(l, m, \lambda^2) = Q(l, m, \lambda^2) + iU(l, m, \lambda^2). \quad (10)$$

An important insight from Burn (1966) is that this is related to the polarized intensity as a function of Faraday depth, F , by

$$P(l, m, \lambda^2) = \int_{-\infty}^{\infty} d\phi F(l, m, \phi, \lambda^2) e^{2i\phi\lambda^2}. \quad (11)$$

We refer to F as the Faraday spectrum. The complex phase term reflects that the position angle of the polarized emission originating at every ϕ location is rotated according to Eq. 7.

Equation 11 is similar to a Fourier transformation, and the ability to invert this relationship is desirable, but two things prevent this. First, it is not possible to sample P for all values of λ^2 . One can only achieve a limited coverage of $\lambda^2 > 0$, and of course cannot measure $\lambda^2 \leq 0$ at all. This problem is addressed by the RM synthesis technique introduced by Brentjens & de Bruyn (2005), where the inversion of Eq. 11 is treated in much the same way as the inversion of Eq. 5. Second, the spectral dependence of F prevents one from inverting Eq. 11. As addressed by Brentjens & de Bruyn (2005), the inversion is possible assuming that the λ^2 and ϕ dependent parts of F are separable, which means that the emission at all ϕ values along a LOS is produced with the same frequency spectrum (up to normalization), i.e.

$$F(l, m, \phi, \lambda^2) = f(l, m, \phi) s(l, m, \lambda^2). \quad (12)$$

In general, of course, this assumption is not valid. Consider the case where two sources lie along the line of sight, each having different emission spectra. By factoring the Faraday spectrum as above we introduce an error into the image of f . Brentjens & de Bruyn (2005) point out that such errors do not effect the Faraday depth of a source, and only have a relatively minor effect on the flux density. In their simulations, using a bandwidth that was 17% of the central frequency, an absolute error of 1 in the spectral index corresponds to an error of less than 5% in flux density. This error should be more pronounced with increased bandwidth, since the change in flux density over the frequency range becomes greater.

Assuming Eq. 12 applies, the dirty Faraday spectrum, f_D , is recovered by

$$f_D(l, m, \phi) = \int_{-\infty}^{\infty} d\lambda^2 S_{\lambda^2}(l, m, \lambda^2) \frac{P(l, m, \lambda^2)}{s(l, m, \lambda^2)} e^{2i\phi\lambda^2} \\ = B_{\phi}(l, m, \phi) * f(l, m, \phi) \quad (13)$$

where $*$ is a convolution with respect to ϕ . The sampling function, S_{λ^2} , is defined to be

$$S_{\lambda^2}(l, m, \lambda^2) = W_{\lambda^2}(l, m, \lambda^2) \sum_i \delta(\lambda^2 - \lambda_i^2), \quad (14)$$

where λ_i^2 are the discrete values of wavelength at which measurements have been made, and W_{λ^2} is a weighting term similar to that in Eq. 4. The so-called RM spread function, or the dirty beam in ϕ -space, B_{ϕ} , is given by

$$B_{\phi}(l, m, \phi) = \int_{-\infty}^{\infty} d\lambda^2 S_{\lambda^2}(l, m, \lambda^2) e^{2i\phi\lambda^2}. \quad (15)$$

Equation 13 shows that, like with aperture synthesis, the product of RM synthesis is a convolution between the true brightness distribution (the Faraday spectrum in this case) and a dirty beam or point-spread function. Again, this dirty beam has structure that extends well beyond the source location. Therefore deconvolution is necessary in order to recover faint sources in the Faraday spectrum and obtain more physically realistic results.

2.3. 2+1D imaging of the Faraday spectrum

RM synthesis is a 1D imaging procedure that is applied to each LOS independently. With aperture synthesis imaging, one obtains the sky brightness distribution across the plane of the sky at a single frequency. Applying RM synthesis to such images requires that the variation of the sky brightness as a function of frequency must be determined at every location in the image plane.

The sampling of the uv -plane, S , varies as a function of frequency. The images obtained from Eq. 5 at different wavelengths will be sensitive to different spatial frequencies and have different resolutions. This complicates measurement of the spectral properties of the sky brightness because changes due to the sampling function can be confused with real variation of the sky brightness distribution. This problem is approximately overcome by applying different weighting and uv -plane tapering schemes to the data such that the uv -coverage is made to be roughly the same at each frequency.

The usual practice of applying the RM synthesis technique to interferometric data involves several steps. We refer to this process as 2+1D imaging:

- Calibrate the visibility data.
- Weight and taper the data such that the resolution of the images is roughly constant with frequency. Some compensation is also needed if there is significant flux missing due to the usual gap in the center of the uv -plane.
- Compile a series of deconvolved Stokes Q and U images at each frequency. Deconvolution is almost always performed using the CLEAN algorithm and the CLEAN model components are convolved with a so-called restoring beam, i.e. a Gaussian profile representing the resolution of the image (determined from the main peak of the dirty beam). The same restoring beam is used for all frequencies.

- Stack the images. Pixel-by-pixel, the polarized intensity as a function of frequency is read from the maps, corrected for spectral variation, s (determined by measuring the spectral index from total intensity maps), and Fourier transformed.
- Perform further ϕ -space deconvolution.

The result is a 3D cube of data representing $f(l, m, \phi)$.

There are two immediate problems with this procedure. First, the necessity to downweight data to approximately match resolutions between images at different frequencies can lead to significant problems. This cannot be done perfectly and any variation in the polarized intensity arising from the differing resolutions produces a shift in the Faraday depth of the emission, thereby introducing systematic errors into the process. Second, Faraday synthesis is performed on maps that have already been processed using non-linear, ad-hoc deconvolution algorithms. Artifacts introduced into the images by said algorithms will be compounded during RM synthesis.

We now introduce a new approach that is a natural extension of the typical synthesis imaging procedures described above, and does not suffer from the problems of the 2+1D imaging technique.

3. Faraday synthesis

In this section we show that it is possible to directly relate the Faraday spectrum to the visibilities of the linearly polarized emission. First, we decompose the Faraday spectrum into parts that relate directly to the intensity distributions of the two Stokes parameters Q and $U(l, m, \lambda^2)$.

Equation 11 can be rewritten as

$$Q + iU = \int_{-\infty}^{\infty} d\phi (F_Q + iF_U) e^{2i\lambda^2\phi} \quad (16)$$

where we have decomposed F into two complex valued terms, F_Q and F_U , that relate directly to the Stokes Q and U brightness distributions, respectively. Therefore,

$$Q = \int_{-\infty}^{\infty} d\phi F_Q e^{2i\lambda^2\phi}, \quad (17)$$

and the same relation holds for Stokes U , as well. We note that because Q and U are real, F_Q and F_U are generally complex and Hermitian in ϕ , i.e. $F_Q^*(l, m, \phi) = F_Q(l, m, -\phi)$. We note further that F_Q and F_U are not the local Stokes Q and U brightnesses at location (l, m, ϕ) , but simply auxiliary variables useful for the formalism. If we assume that F can be factored as in Eq. 12, then

$$Q(l, m, \lambda^2) = s(l, m, \lambda^2) \int_{-\infty}^{\infty} d\phi f_Q(l, m, \phi) e^{2i\lambda^2\phi} \\ = s(l, m, \lambda^2) q(l, m, \lambda^2). \quad (18)$$

We will now work only with Stokes Q , but we note that the following expressions also hold for Stokes U and F_U . Following Eq. 3, the measured Stokes Q visibility, \widehat{V}_Q , is

$$\widehat{V}_Q = S \iint_{-\infty}^{\infty} dl dm A Q e^{-2\pi i(ul+vm)} \\ = S \iint_{-\infty}^{\infty} dl dm A s q e^{-2\pi i(ul+vm)} \quad (19)$$

where S is defined as in Eq. 4. The relationship between the Stokes visibilities and the correlator output depends on the type of antenna feeds that are used. For linearly polarized feeds

$$\begin{aligned} V_Q &= V_{XX} - V_{YY} \\ V_U &= V_{XY} + V_{YX}, \end{aligned} \quad (20)$$

where V_{XX} , V_{YY} , etc. represent the visibilities from cross-correlations of feeds having X or Y perpendicularly oriented dipoles. For circularly polarized antenna feeds,

$$\begin{aligned} V_Q &= V_{RL} + V_{LR} \\ V_U &= -i(V_{RL} - V_{LR}), \end{aligned} \quad (21)$$

where V_{RL} and V_{LR} represent the visibilities from the cross-correlation between feeds of right and left, or left and right circular polarization, respectively.

We now define $a(l, m, \phi)$ and $\sigma(l, m, \phi)$ to be the representations of the primary beam and spectral dependence of the Faraday spectrum in Faraday space, respectively, i.e.

$$A(l, m, \lambda^2) = \int_{-\infty}^{\infty} d\phi a(l, m, \phi) e^{2i\lambda^2\phi} \quad (22)$$

and

$$s(l, m, \lambda^2) = \int_{-\infty}^{\infty} d\phi \sigma(l, m, \phi) e^{2i\lambda^2\phi}. \quad (23)$$

Using the convolution theorem in between ϕ and λ^2 , we combine the expressions above to find that the Faraday spectrum is related to the measured visibilities by

$$\widehat{V_Q} = S \iiint_{-\infty}^{\infty} dl dm d\phi \left(a *_{\phi} \sigma *_{\phi} f_Q \right) e^{-2\pi i \left(ul + vm - \frac{\lambda^2}{\pi} \phi \right)}, \quad (24)$$

where $*_{\phi}$ denotes a 1D convolution with respect to ϕ along each LOS. This expression can be inverted to give the dirty image for the Faraday spectrum

$$\begin{aligned} \left(a *_{\phi} \sigma *_{\phi} f_Q \right)_D &= \iiint_{-\infty}^{\infty} du dv d\lambda^2 \widehat{V_Q} e^{2\pi i \left(ul + vm - \frac{\lambda^2}{\pi} \phi \right)} \\ &= B * (a *_{\phi} \sigma *_{\phi} f_Q), \end{aligned} \quad (25)$$

where

$$B(l, m, \phi) = \iiint_{-\infty}^{\infty} du dv d\lambda^2 S(u, v, \lambda^2) e^{2\pi i \left(ul + vm - \frac{\lambda^2}{\pi} \phi \right)}, \quad (26)$$

and $*$ is a full 3D convolution in l , m , and ϕ . We note that the 3D and 1D convolution operations do not commute. We can see that by inverting the visibilities, we do not directly recover the Faraday spectrum, but rather the Faraday spectrum convolved with B in 3D, and σ , and a in 1D along each LOS. In order to recover f_Q , one must first perform a 3D deconvolution using the CLEAN algorithm, for example. After the 3D deconvolution, deconvolution of a and σ can be achieved by performing a 1D inverse Fourier transform into λ^2 -space along each LOS, dividing by A and s , and then Fourier transforming back into ϕ -space.

The beam pattern A is usually known to high precision and is often represented by an analytic function parameterized in l , m ,

and λ^2 . In addition, a map of s will be required. This can be obtained by measuring the spectral variation of total intensity maps along each LOS. In many circumstances, it may be sufficient to simply assume that s is independent of l and m , since the errors introduced by using the wrong form for s are often quite small as previously discussed.

Imaging software that implements the 3D inversion given by Eq. 25 would avoid the complications described for traditional, 2+1D imaging. We eliminate the need to match uv -coverage at all frequencies because the 3D dirty beam, B , is constructed from the full 3D sampling function. We also avoid the possibility of compounding errors through the process of deconvolving images that have already once been deconvolved. As a result, the fidelity of the images produced using Faraday synthesis should be improved over those made using the 2+1D technique. In principle, the 3D approach will result in images that have higher dynamic range than those obtained using the 2+1D approach because we are able to use all data across the full bandwidth during imaging and deconvolution.

We have presented a simplified description of the Faraday spectrum measurement process above. This will already work quite well in many circumstances, notably for narrow-field observations without significant direction dependent effects, but in general additional steps will be required. For instance, when the w -term in Eq. 1 can not be ignored, the w -projection algorithm of Cornwell et al. (2005) has been shown to be very effective in reducing imaging errors. This algorithm makes use of the fact that the multiplication of the w -term and the sky brightness in the image plane is a convolution in visibility space. The w -dependent visibilities are projected onto the $w = 0$ plane using the convolution kernel, and then a 2D Fourier transform can be used to recover the sky brightness distribution. This algorithm can be applied to the case of Faraday synthesis without modification. It should also be possible to apply the A-projection algorithm of Bhatnagar et al. (2008), which corrects for direction-dependent beam effects in a manner similar to the w -projection algorithm.

4. Proof of concept implementation

To compare the 3D approach to polarization imaging with the traditional 2+1D approach, we have implemented a proof of concept 3D imaging and deconvolution software package called `fsimager`. For deconvolution, a 3D CLEAN algorithm has been implemented because it is by far the most common deconvolution method used in radio astronomical imaging, and with it we can make the most direct comparison between the two techniques.

The CLEAN algorithm (Högbom 1974) is a non-linear, iterative deconvolution routine. The algorithm makes the implicit assumption that the sky is composed of point sources distributed throughout a mostly blank field. Over the last decades, it has been shown to work quite well even for fields that do not strictly meet this criterion. In brief, the procedure calls for iteratively building up a model of the sky by locating the peak of the image, adding a point source to the model at the location of the peak and with some fraction of its strength, and subtracting from the image the new model point convolved with the dirty beam. This is repeated until one can add no further flux to the sky model, i.e. when one is CLEANing the noise. A complete description of the 3D CLEAN algorithm that we have implemented is given in Appendix A.

Radio astronomical imaging relies heavily on Fourier transforms, and because the number of pixels in the 3D image of the Faraday spectrum is large, fast Fourier transforms (FFTs)

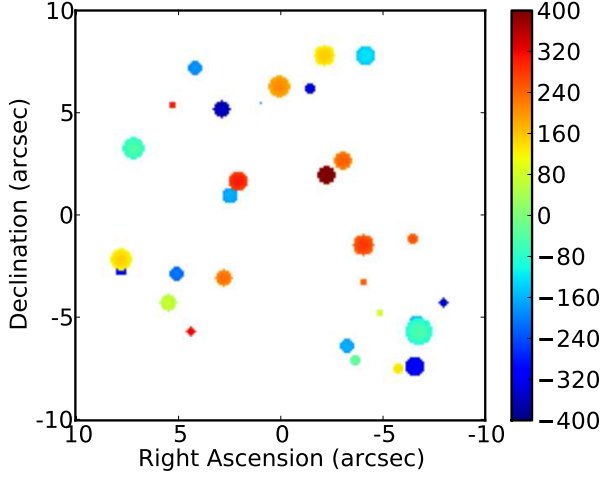


Figure 1. The sky-plane distribution of the 30 point sources in the model. The color scale indicates the Faraday depth, given in rad/m^2 , and the area of each circle is proportional to the log of the polarized flux.

are required for computational feasibility. Visibility data is not collected on a regularly spaced grid in (u, v, λ^2) -space. To use the FFT algorithm, the data must be interpolated onto regularly spaced grid points prior to processing. For this we employ a well known procedure known as *gridding* that is used extensively in aperture synthesis imaging as well as in medical imaging. Specifically, we have implemented an algorithm described by Beatty et al. (2005). In brief, we convolve the data with a Kaiser-Bessel window function (KBWF) and sample the result on a regular grid in (u, v, λ^2) -space. After this procedure, one is able to perform a FFT with the same result achieved by using a discrete Fourier transformation, to within an arbitrarily small accuracy. The attenuation of the image plane caused by the convolution with the KBWF is corrected for by dividing the dirty image by the Fourier transform of the KBWF. There are two parameters in the gridding procedure that affect precision at the cost of computational time. One parameter describes the extent of the KBWF in visibility space. The second parameter, the so-called oversampling ratio, is the factor by which the image plane should be enlarged in order to mitigate problems that occur at the edges of the image. With the parameters that we have chosen, for the KBWF to extend over 6 pixels in each of u , v , and λ^2 , and an oversampling ratio of 1.5, the results are the same as one would obtain with a discrete Fourier transformation (DFT) to within one part in 10,000. In fact the dynamic range is only so limited near the edges of the image where the effects of the convolution with the KBWF are the most dramatic. Away from the edges of the image, the dynamic range is much higher. Overall, the dynamic range can easily be improved by changing the gridding parameters, but this is done at the expense of increased processing and memory requirements.

We have implemented *fsimager* in Python, thus allowing for rapid development and testing, but resulting in overall poor performance. To improve the situation, we have optimized some sub-functions (particularly the gridding routines) using Cython and in the end the code performs admirably. We can load a 1GB data set, grid, image, and CLEAN a 16 megapixel image using 500 iterations in about 15 minutes on our 2.4 GHz Core i5 development machine with 8 GB of RAM. The most major limitation of the current version of the software is that all of the data re-

sides in memory and this limits the size of the images that can be produced to the amount of memory that is available on the machine. We have run all tests on a computer having 64 GB of memory, and are limited to producing image cubes that are 400 megapixels in size (about 750 pixels per side). This may be sufficient for imaging small fields of view, but when imaging data from a wide-field, high resolution instrument (e.g. LOFAR), this is inadequate.

5. Tests

To compare the 3D and 2+1D approaches, we have produced a mock observation of a set of polarized point sources. In this section, we first describe the mock data, and then we describe the method that we use to image the data using the 2+1D method. Finally, we compare the Faraday spectra produced by each technique.

5.1. Mock data

We create a model Faraday spectrum consisting of a number of polarized point sources. We choose this simple scenario, though it may not be the most physically motivated one, because it is what the CLEAN algorithm was designed for. In these tests we simply want to study the differences between the two imaging frameworks, not the merits of the specific imaging algorithm being used. By choosing a model for which CLEAN is ideally suited, we can concentrate on the fundamental differences between the 3D and 2+1D techniques without being concerned about artifacts introduced by using an inappropriate deconvolution algorithm.

No spectral variation is included in the model Faraday spectrum, i.e. F is only a function of l , m , and ϕ . Also, we assume that our field of view is small relative to the size of the primary beam, and therefore that there is no variation in beam strength throughout the image plane.

Our model includes 30 polarized point sources, randomly distributed throughout the volume. The distribution of model sources is shown in Fig. 1. The color indicates the Faraday depth, where the scale shown on the right is in rad/m^2 . The size of the points in the plot indicates the relative linearly polarized flux of the source. The fluxes range from 0.06 Jy to 64 Jy, thus allowing for a comparison across a wide range of signal to noise ratios.

We "observe" the model by taking its Fourier transformation and then filling the data column of a custom made MeasurementSet file. The MeasurementSet is created using the *makems* script that is part of the LOFAR software package. This script creates a MeasurementSet with user specified time and frequency coverages, and computes uv -coordinates using the antenna table from a pre-existing MeasurementSet. We use the antenna table from a VLA A-array observation. The frequency coverage is specified such that our mock data spans the range from 1-4 GHz, with 64 frequency channels distributed throughout the range. The total time covered by the observation is 1 hour, with a 60 second step size between measurements to reduce the file size.

Gaussian white noise, having a standard deviation of 10 Jy, is added to real and imaginary parts of the Stokes Q and U visibilities separately. The MeasurementSet file is then either read directly into the 3D software for gridding, imaging, and deconvolution, or into the CASA software package for traditional imaging.

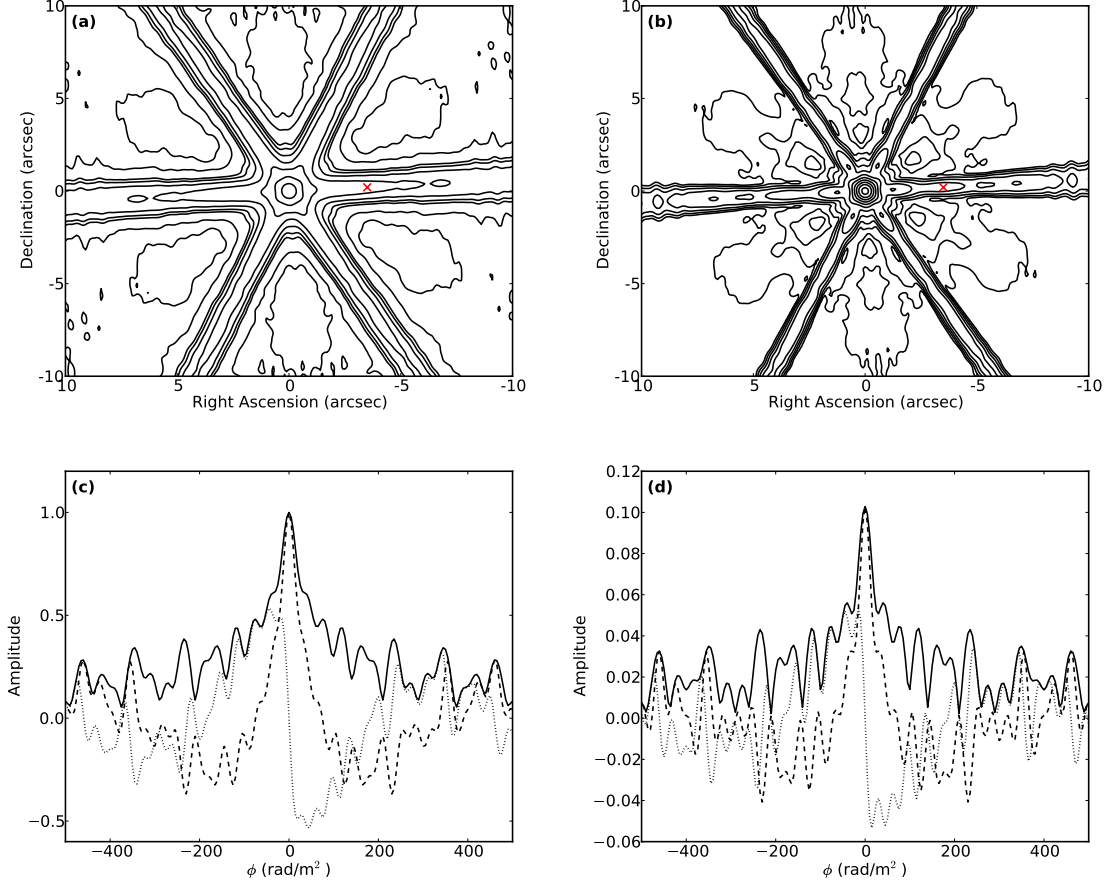


Figure 2. The 3D dirty beam. a) A 2D slice through the cube in the sky plane at $\phi = 0 \text{ rad/m}^2$. The magnitude of the complex valued beam is shown. The bottom contour begins at 0.01 and there is a step of 2 between levels. Negative contours are marked with dashed lines. b) Another view of the beam, now at $\phi = 50 \text{ rad/m}^2$. c) A 1D profile through the center of the sky plane at $(0'', 0'')$. d) A 1D profile along the LOS at $(-3.5'', 0.2'')$, marked with an “x” in Figs. a and b. The magnitude of the profiles are shown as a solid line, while dashed and dotted lines indicate the real and imaginary parts, respectively.

5.2. 2+1D imaging procedure

The 2+1D imaging procedure is conducted by first loading the mock observation data into CASA where it is then written to UVFITS format using the task *exportuvfits*. This data is loaded into AIPS and each channel is independently deconvolved using the IMAGR task.

In an attempt to match the uv -coverage at the different frequencies, weighting and tapering schemes are employed such that the output image resolutions are roughly equal. At 1 GHz, the maximum uv -spacing is approximately $125 \text{ k}\lambda$ and the main peak of the synthesized beam has a FWHM of approximately $1.7''$. At 4 GHz, the maximum uv -spacing is $490 \text{ k}\lambda$ and the FWHM of the main peak of the synthesized beam is approximately $0.6''$. Without tapering, we will observe large variations in the sources as a function of frequency simply due to the dramatic change in resolution.

All but the lowest frequency observations are tapered in the uv -plane with a Gaussian profile having a FWHM of $150 \text{ k}\lambda$. Robust weighting is used, and the weighting parameter is chosen such that FWHM of the main peak of the synthesized beam is approximately $1.2''$. After deconvolution in the image plane, all images are restored using a Gaussian profile that has a FWHM of $1.2''$. We have tried a few other weighting and tapering schemes so to achieve even lower resolutions, down to a FWHM of $1.7''$.

The different schemes made no significant difference in the results apart from the resolution of the final image cube.

RM synthesis is performed along each LOS using our own RM synthesis software that is currently in use on the LOFAR compute cluster for commissioning. This software implements the RMCLEAN algorithm described by Heald et al. (2009).

5.3. Results

The result of either the 2+1D or 3D imaging procedures is a 3D reconstruction of $F(l, m, \phi)$. In Fig. 3 we show a side-by-side comparison of the reconstructed Faraday spectra for a few selected Faraday depths. In each row, the left panel in each image shows the model image, the middle panel shows the 2+1D reconstruction and the Faraday synthesis reconstruction is shown on the right.

The most obvious difference between the two images is the resolutions. In the 3D reconstruction, we use a natural weighting scheme. The FWHM of the main peak of the 3D dirty beam is roughly $0.8''$, $0.8''$, and 40 rad/m^2 in R.A., Dec., and ϕ , respectively. Selected image plane slices, and LOS profiles of the Faraday synthesis derived 3D dirty beam are shown in Fig. 2. For the 2+1D imaging, we have in some sense chosen the resolution in the sky plane to be $1.2''$ by selecting a particular weighting

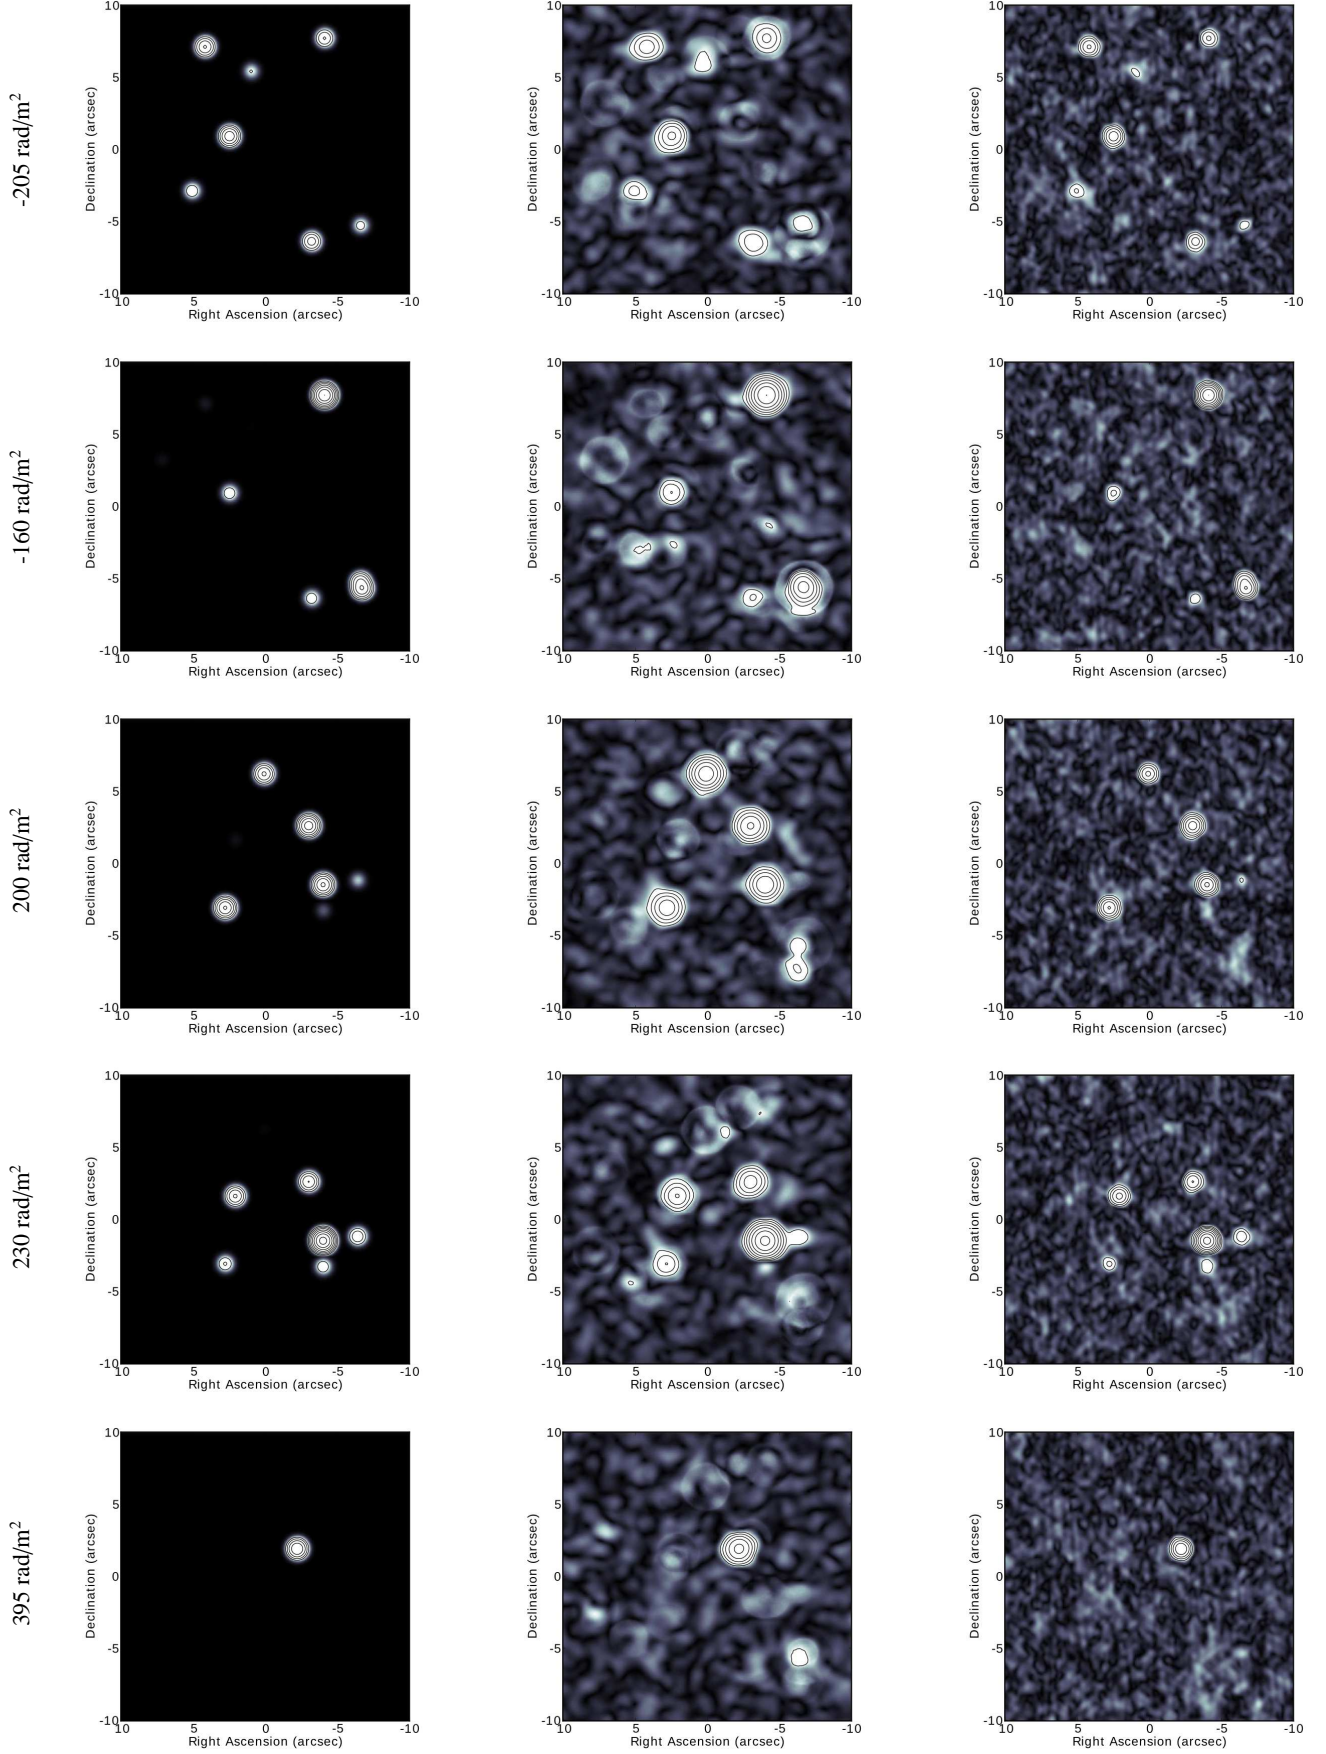


Figure 3. A comparison between the input model and the two reconstructions at several ϕ values. In each frame, the greyscale varies linearly from 0 – 50 mJy/beam , and the contours begin at 50 mJy/beam with a factor of 2 between each level. Left: The model image convolved with an $(0.8'' \times 0.8'' \times 40 \text{ rad/m}^2)$ Gaussian. Middle: The 2+1D reconstruction. Right: The *fsimager* reconstruction. Top to bottom: $\phi = -205, -160, 200, 230$, and 395 rad/m^2 .

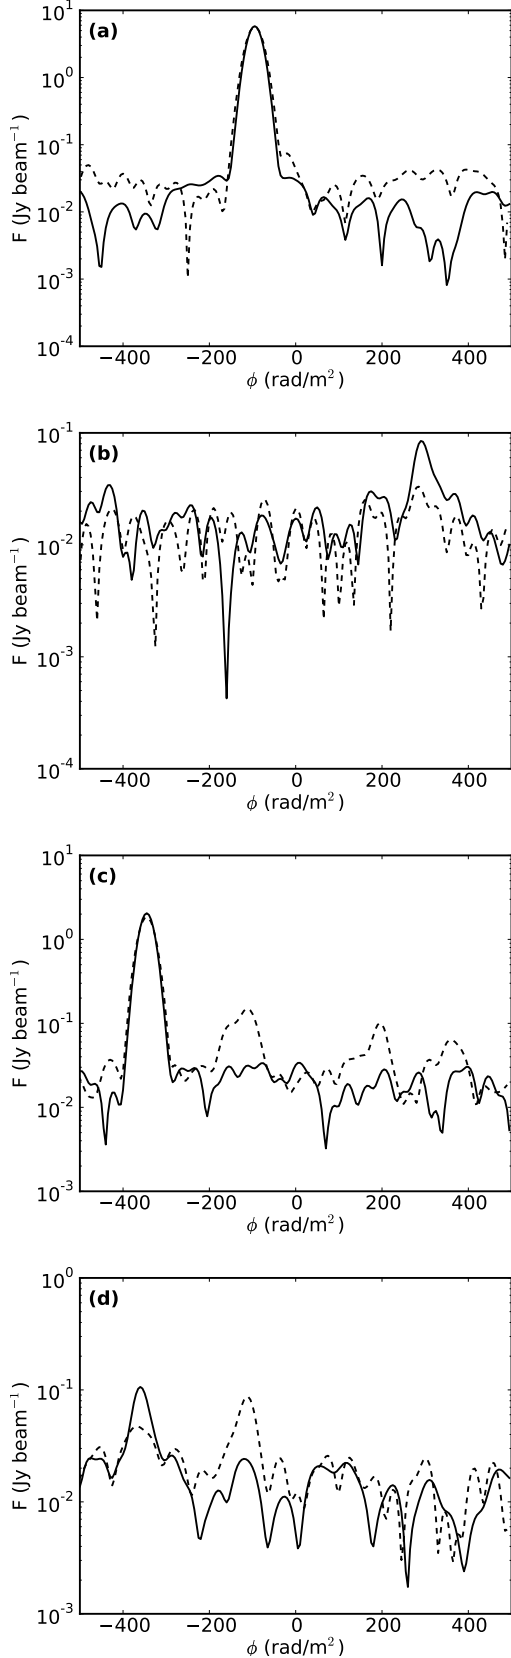


Figure 4. A few LOS profiles through the reconstructed Faraday spectra. Note that the brightness scale is logarithmic. In each, the solid and dashed lines indicate the 3D and 2+1D reconstructions, respectively. The profiles are taken along the following image plane coordinates: a) (7.2'', 3.2''), b) (5.3'', 5.3''), c) (-6.5'', -7.4''), d) (-7.9'', -4.3'').

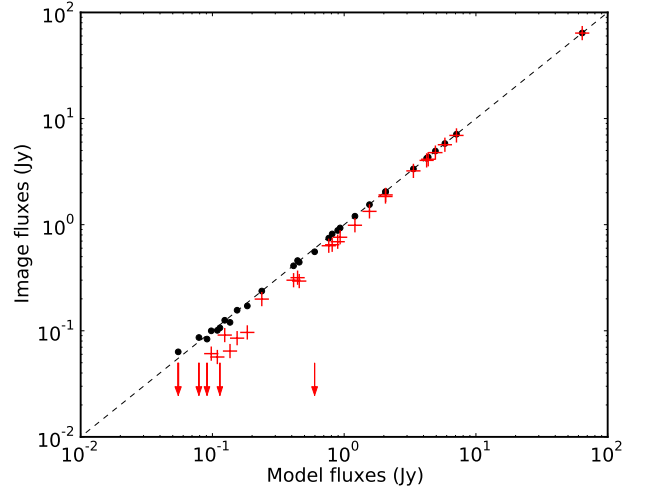


Figure 5. Reconstructed source fluxes compared to those of the input model. Dark circles denote the fsmager reconstruction while light crosses denote the 2+1D reconstruction. Arrows indicate upper levels of reconstructed fluxes that were not located by our search algorithm.

scheme, but our choice is limited by the resolution of the lowest frequency part of the data. We are therefore able to achieve a higher resolution with the Faraday synthesis technique, without the need for tapering the long baseline data.

The LOS profile through the center of the 3D dirty beam, shown in Fig. 2c, is the same as the usual RM synthesis dirty beam given in Eq. 15. However, the off-center profile, shown in Fig. 2c, is not simply a scaled version of the usual RM synthesis dirty beam. The structure of the two profiles is quite different. This is because the sidelobes of B_{sky} change as a function of frequency, which leads to structure in ϕ -space.

We find that the noise level in the 2+1D image is higher than that of the 3D image. We measure the noise levels in the two reconstructions by computing the root mean square (RMS) pixel values in several regions of the image cubes where no sources are located. In the 3D image cube the noise is 6.06 mJy/beam. The RMS is 20% higher in the 2+1D image cube, measuring 7.33 mJy/beam.

Both imaging techniques are able to successfully reconstruct the stronger sources in the model image quite well. The sources above 1 Jy have all been located correctly, and the recovered fluxes are within a few percent of the model fluxes in each case. In the reconstruction at $\phi = -205 \text{ rad/m}^2$, shown in the first row of Fig. 3, both the 2+1D and 3D reconstructions roughly match the input model. Even the weak sources are detected. This represents one of the better slices of the 2+1D reconstruction, and yet here we can already begin to see problems. There are some hints of artifacts in the traditionally made image that are not present in the Faraday synthesis result. The artifacts appear prominently in most other frames of the reconstruction.

Overall, the blank sky in the Faraday synthesis reconstruction is noise-like throughout the image volume, while in contrast, we find many artifacts present in the 2+1D reconstruction. The artifacts often appear as circular outlines of sources at different Faraday depths, like the feature at $\sim(7'', 3'')$ in the image at $\phi = -160 \text{ rad/m}^2$. Such features are present in each of our example image slices above. These circular features are quite obviously artifacts, but in some cases, for example at $\sim(-6'', -7'')$

in the image at $\phi = 200 \text{ rad/m}^2$, false features are quite strong and not as obviously artificial. One can see a weaker analog of this particular artifact in the 3D reconstruction as well, but most other artifacts in the 2+1D images have no counterpart in the 3D reconstruction. These erroneous features act to confuse the detection of weaker sources in the field.

There are artifacts that appear along the LOS profiles of the 2+1D reconstruction as well. In Fig. 4, we show a few LOS profiles through each image cube. Figure 4a shows a profile through a relatively strong and isolated source at R.A.=7.2", and Dec.=3.2" (c.f. Fig. 1). This source is well reconstructed by both methods. In the other example profiles, problems with the 2+1D reconstruction become apparent. The profile shown in Fig. 4b is again through a rather isolated but weaker source, at R.A.=5.3", Dec.=5.3". Here we find that the actual source is not prominent compared with the noise in the 2+1D profile, but is easily detected in the 3D reconstruction.

The other profiles are through more crowded regions of the image cube, where contamination from nearby sources becomes problematic. The profile in Fig. 4c, along R.A.=−6.5" and Dec.=−7.4", passes through another relatively strong source. While the source itself is well reconstructed, three other features due to nearby sources appear around $\phi = -150, 200$, and 350 rad/m^2 in the 2+1D reconstruction. These are not present in the 3D case. Figure 4d shows a profile through a relatively weak source at R.A.=−7.9", and Dec.=−4.3". The source is easily distinguishable in the 3D profile, but in the 2+1D profile it is completely overshadowed by the artifact attributed to the source at R.A.=−6.7", and Dec.=−5.7", almost 2" away.

These problems occur in part due to the lower resolution of the 2+1D reconstruction. They are also due to the fact that each channel in the data is imaged and deconvolved (in 2D) separately, and therefore with a more limited sensitivity. As a result, residual artifacts remain in the individual images, and these become apparent with the increased sensitivity provided by combining data over the full frequency range. These residuals are then processed during RM synthesis, and can become quite problematic as we can clearly see in our examples.

In each reconstruction, even the weakest model source should be present at the 10σ level in the 3D image and at 8σ in the 2+1D case. To locate point sources in each image, we have a simple routine for locating local maxima that scans through every pixel and checks whether it is larger than all neighboring pixels. If so, the location and brightness is recorded. We search through each image cube in this way to find all local maxima above the 50 mJy/beam level. Ideally we should expect to find 30 points corresponding to the model sources. In the 3D reconstruction, we find 32 such locations. All input model sources are located along with two false detections. In the 2+1D case we locate 147 sources, with 5 input model sources not detected. A few of the missing sources are not detected because they are simply not resolved from a nearby stronger source. The others may be present below the 50 mJy/beam cutoff, but at this point they are completely indistinguishable from the artificial sources.

The artificial sources in the 2+1D reconstruction are not only much more numerous, but also brighter. The brightest artifact in the 2+1D image cube is 0.113 Jy/beam , compared to 0.062 Jy/beam in the 3D image cube.

Figure 5 compares the fluxes of the sources found in the two reconstructions to those of the input model. These have been corrected for the Ricean bias effect described in Wardle & Kronberg (1974). The 3D reconstruction agrees remarkably well with the model fluxes across the full range of source strengths. This sky model is ideally suited for the CLEAN algorithm, so this is as

should be expected. The stronger sources are also recovered nicely in the 2+1D reconstruction, but the strengths of weaker sources are systematically lower than the input model fluxes.

6. Discussion and conclusions

We have described a new approach to imaging linearly polarized visibility data that we call Faraday synthesis. With this approach, one directly reconstructs the Faraday spectrum, or the polarized intensity as a function of Faraday depth, from the visibility data. This takes the place of the usual approach of first performing aperture synthesis imaging on the visibility data at each frequency, then performing RM synthesis along each line of sight independently.

These two approaches would be equivalent if deconvolution were not required. With Faraday synthesis, the deconvolution is done in one step using the entirety of the broad-band data. In contrast, the traditional approach requires deconvolving the images individually at each frequency. The sensitivity in the narrow-band images is significantly limited, and residual artifacts remain in these images that limit the dynamic range and image fidelity achieved during RM synthesis. Moreover, another deconvolution procedure is necessary to account for the point spread function due to the incomplete sampling of the wavelength space. Artifacts introduced by the first deconvolution algorithm will be compounded during this procedure, further reducing image fidelity.

Indeed, we found in our proof-of-concept testing that artifacts were significantly higher in the traditional imaging method than with Faraday synthesis. The noise was roughly 20% lower when using the Faraday synthesis technique, and the strongest artifact was about half as bright.

We found that one is able to achieve a better resolution in the final image using the new approach. The main lobe of the 3D dirty beam was 30% smaller in the sky plane than that of the traditional method. With the 2+1D method, stacking of the images at each frequency requires tapering the visibility data such that the higher frequency images have the same resolution as those at the lower frequencies. We find that the inaccuracies inherent in the stacking process are a significant source of artifacts in the traditional RM synthesis technique. Furthermore, this procedure requires severely down-weighting a large portion of the available data, again limiting sensitivity. With the Faraday synthesis approach, no such down-weighting is required.

The Faraday synthesis approach of working directly between visibility and Faraday space is a much better foundation on which to build new image reconstruction algorithms because with it one is able to accurately describe the instrument response. The effects of the intermediate, non-linear deconvolution procedure can not be easily understood or modeled. Many signal inference techniques, like those derived using Information Field Theory (Enßlin et al. 2009; Enßlin & Weig 2010), depend on a complete description of the instrument response and would need to be built on the framework of Faraday synthesis.

While the CLEAN algorithm has worked quite well in radio astronomy for decades, the implicit assumption of a sky sparsely populated by point sources is not well suited for the kinds of diffuse polarized signals that one finds, for example, in the polarized emission from the Milky Way. Novel image reconstruction algorithms built using more appropriate constraints or assumptions are likely justified. Such algorithms could make use of statistical correlations inferred from the data, similar to the extended critical filter algorithm developed by Oppermann et al. (2011).

Acknowledgements. This research was performed in the framework of the DFG Forschergruppe 1254 Magnetisation of Interstellar and Intergalactic Media: The Prospects of Low-Frequency Radio Observations. We thank Henrik Junklewitz, Niels Oppermann, Marco Selig, Maximilian Uhlig, George Heald, and Ger de Bruyn for many helpful discussions. We thank Ger van Diepen at ASTRON for the *makems* software that we use to produce our mock data.

References

- Andrecut, M., Stil, J. M., & Taylor, A. R. 2011, arXiv:1111.4167
- Beatty, P. J., Nishimura, D. G., & Pauly, J. M. 2005, IEEE Transactions on Medical Imaging, 24, 799
- Bell, M. R., Junklewitz, H., & Enßlin, T. A. 2011, Astronomy & Astrophysics, 535, A85
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, Astronomy & Astrophysics, 487, 419
- Brentjens, M. A. 2011, Astronomy & Astrophysics, 526, A9
- Brentjens, M. A. & de Bruyn, A. G. 2005, Astronomy and Astrophysics, 441, 1217
- Burn, B. J. 1966, Monthly Notices of the Royal Astronomical Society, 133, 67
- Clark, B. G. 1980, Astronomy & Astrophysics, 89, 377
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2005, in ASP Conference Series, Vol. 347, Astronomical Data Analysis Software and Systems XIV
- de Bruyn, A. G. & Brentjens, M. A. 2005, Astronomy and Astrophysics, 441, 931
- Enßlin, T. A., Frommert, M., & Kitaura, F. S. 2009, Physical Review D, 80, 105005
- Enßlin, T. A. & Weig, C. 2010, Phys. Rev. E, 82, 051112
- Frick, P., Sokoloff, D., Stepanov, R., & Beck, R. 2010, Monthly Notices of the Royal Astronomical Society, 401, L24
- Gaensler, B. M., Landecker, T. L., Taylor, A. R., & POSSUM Collaboration. 2010, in Bulletin of the American Astronomical Society, Vol. 42, American Astronomical Society Meeting Abstracts #215, 515
- Heald, G., Braun, R., & Edmonds, R. 2009, Astronomy and Astrophysics, 503, 409
- Högbom, J. A. 1974, Astronomy and Astrophysics Supplement, 15, 417
- Li, F., Brown, S., Cornwell, T. J., & de Hoog, F. 2011, Astronomy & Astrophysics, 531, A126
- Oppermann, N., Robbers, G., & Enßlin, T. A. 2011, Physical Review E, 84, 041118
- Pen, U.-L., Chang, T.-C., Hirata, C. M., et al. 2009, Monthly Notices of the Royal Astronomical Society, 399, 181
- Pizzo, R. F., de Bruyn, A. G., Bernardi, G., & Brentjens, M. A. 2010, 1008.3530
- Schwab, F. R. 1984, Astronomical Journal, 89, 1076
- Thompson, A. R., Moran, J. M., & Swenson, Jr., G. W. 2001, Interferometry and Synthesis in Radio Astronomy, 2nd Edition, ed. Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr.
- Wardle, J. F. C. & Kronberg, P. P. 1974, Astrophysical Journal, 194, 249
- Wolleben, M., Fletcher, A., Landecker, T. L., et al. 2010, 1011.0341
- Wolleben, M., Landecker, T. L., Carretti, E., et al. 2009, in IAU Symposium, Vol. 259, IAU Symposium, 89–90

Appendix A: 3D RMCLEAN

In our proof of concept software we have implemented a variant of the CLEAN routine that was introduced by Clark (1980). In this variant the model is populated by searching for peaks in image space, but the model subtraction is performed in uv -space.

We first produce the dirty images of F_{QD} and F_{UD} . These are combined into a single complex image according to

$$\begin{aligned}\Re(F_D) &= \Re(F_{QD}) - \Im(F_{UD}) \\ \Im(F_D) &= \Re(F_{UD}) + \Im(F_{QD}),\end{aligned}\quad (\text{A.1})$$

where \Re and \Im are operators that select the real and imaginary parts of the image, respectively. We proceed with the CLEAN procedure using the complex valued F_D , and the associated visibilities.

The algorithm is performed in two parts, the so-called major and minor cycles. In the minor cycle, new model sources are located in image space. The subtraction of the model from the visibility data is performed during the major cycle. What remains

after the model is subtracted from the visibilities are called the residual visibilities, V_R .

Before starting the procedure, we extract a patch from the dirty beam, B_{patch} , for use during the minor cycle. We record the value $\beta = \max(|B - B_{\text{patch}}|)$, where $|\cdot|$ indicates that we take the magnitude of the complex valued map.

The major cycle includes:

1. Invert V_R to create a residual image of the complex Faraday spectrum, F_R . Find $F_{\text{lim}} = \max(|F_R|)$.
2. Start the minor cycle using F_R and populate a minor cycle sky model, M_{minor} . The minor cycle is described below.
3. Upon completion of the minor cycle, inverse Fourier transform the minor cycle sky model into visibility space, V_M .
4. Subtract the model from the residual visibility data, $V_R = V_R - S V_M$.
5. Add M_{minor} to the complete sky model, M_{major} .
6. Repeat steps 1–6 until some user defined number of iterations has been done, or F_{lim} is below a user defined cutoff.
7. Invert V_R to produce a final residual image. Add to this M_{major} convolved with a Gaussian restoring beam.

The minor cycle proceeds as follows:

1. Find $(l_m, m_m, \phi_m) = \text{argmax}(|F_R|)$.
2. Add a point source to V_M at (l_m, m_m, ϕ_m) with a flux $F_M = g F_R(l_m, m_m, \phi_m)$, where g is a user defined gain factor (between zero and one).
3. Subtract $F_M B_{\text{patch}}$, centered on (l_m, m_m, ϕ_m) , from F_R .
4. Continue as long as $|F_M| > \beta F_{\text{lim}}(1 + \frac{1}{N})$ where N is the total number of minor cycle iterations that have been completed.

We adopt the minor cycle stop condition suggested in Clark (1980). This condition reflects the fact that during the minor cycle we only subtract a small patch of the dirty beam from the image and therefore any effects outside of this patch will not be removed from the image. During the minor cycle, we CLEAN only as deeply as the maximum contribution of the brightest source to the residual image that is *not* removed by subtracting the beam patch. Actually, the $(1 + \frac{1}{N})$ term makes this condition even more strict early on in the procedure.